2010 Rocky Mountain IPv6 Summit

Danny McPherson dmcpherson@verisign.com

Mike Hollyman mhollyman@arbor.net

IPV6 ON THE INTERNET

EMPIRICAL OBSERVATIONS

Overview

- Various stats from infrastructure-enabling elements
  - Allocations
  - Routing
  - DNS
  - Traffic
  - End Systems
- Stats are just stats, use them to gauge adoption, plan roll-out, optimize engineering, and fix brokenness
- Perhaps the stats that matter most are end-end service-level transactions – illustrates all the network glue is working
- As IPv4 free pool exhausts all bets are off…

Conceptual Model

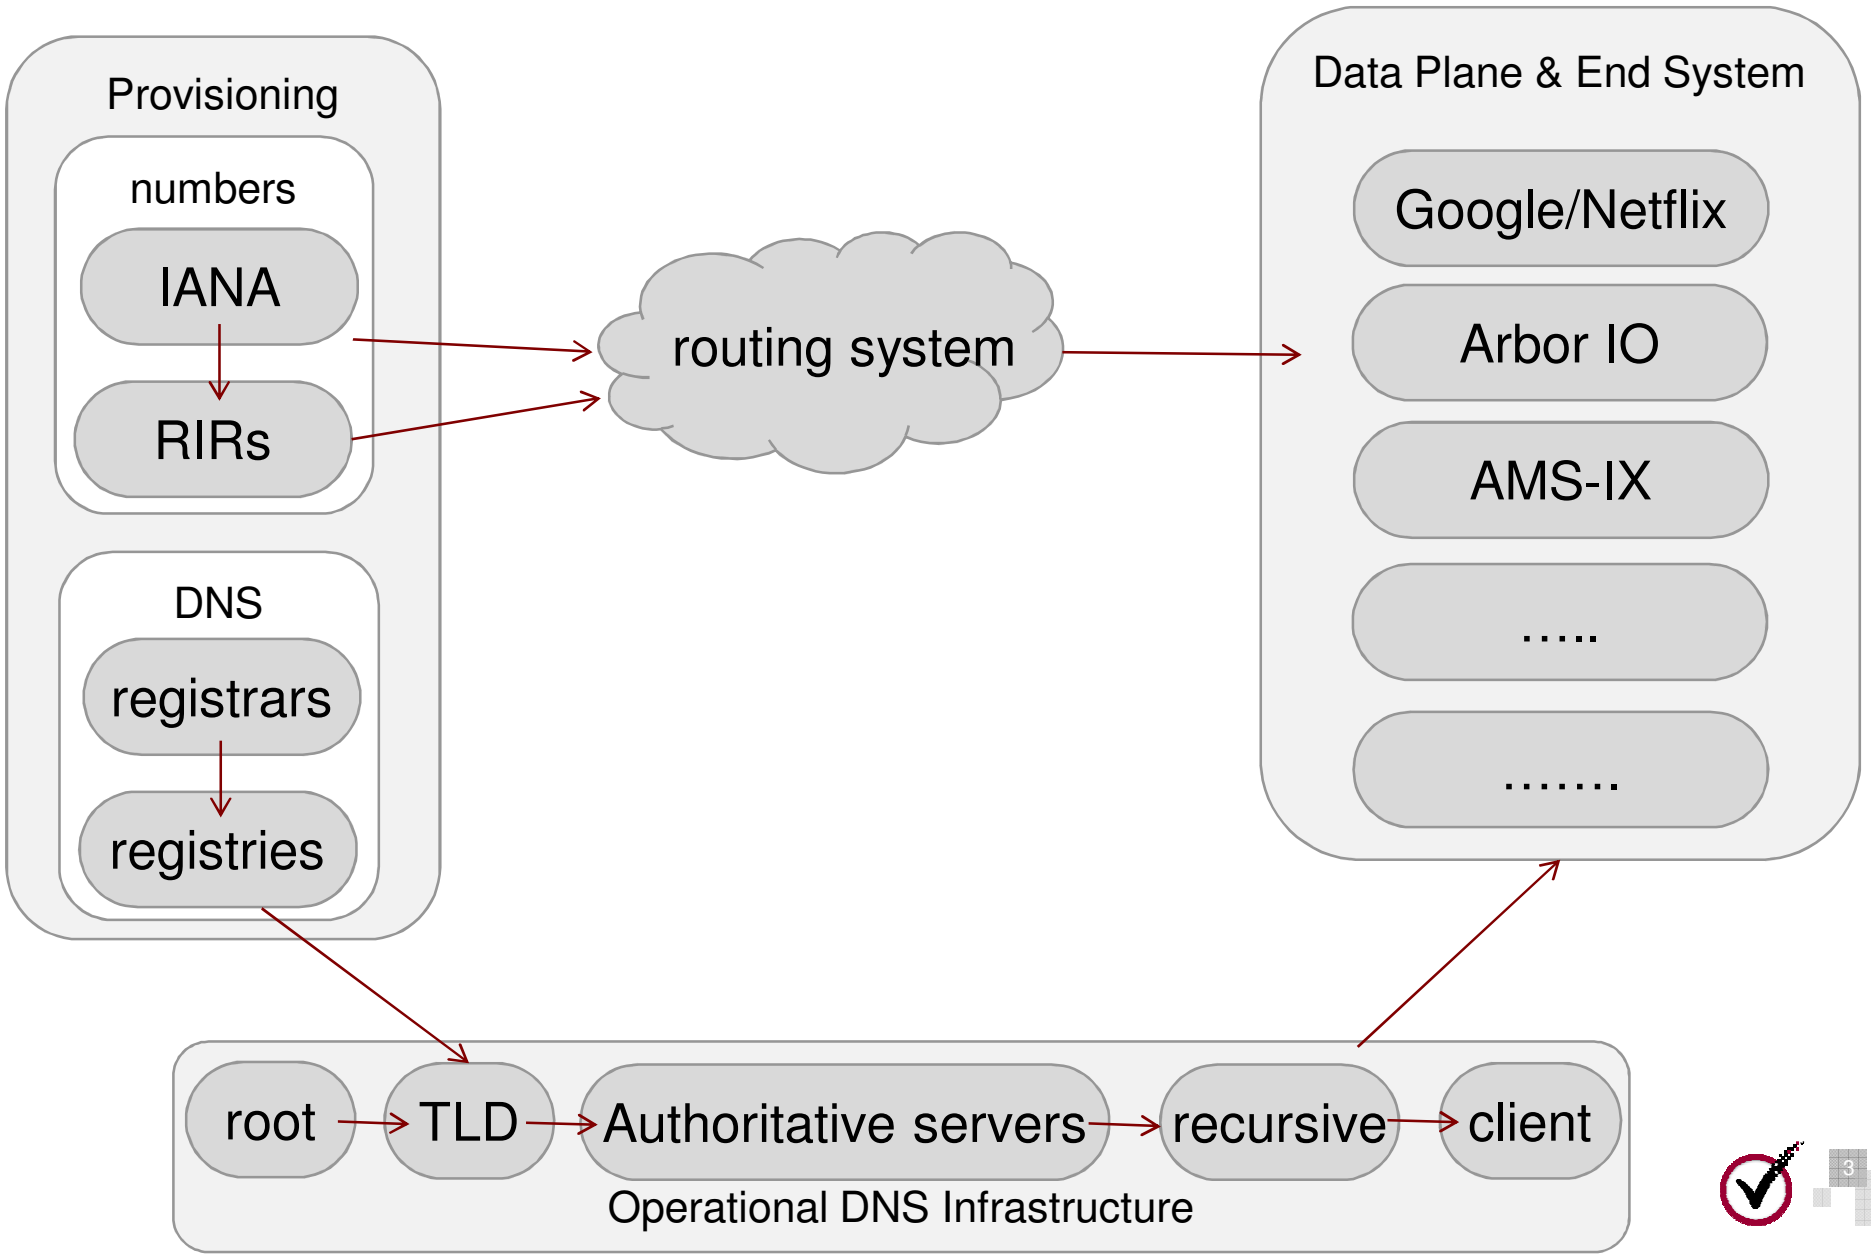

IPv4 Free Pool Exhaustion

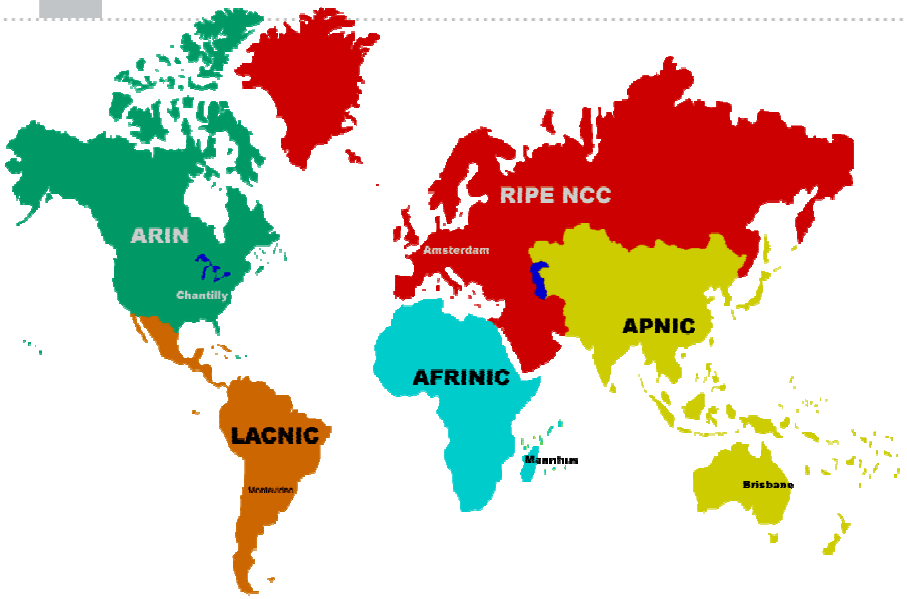

- IANA exhausts IPv4 free pool 30 JUL 2011
- RIRs exhaust free pool 13 MAR 2012

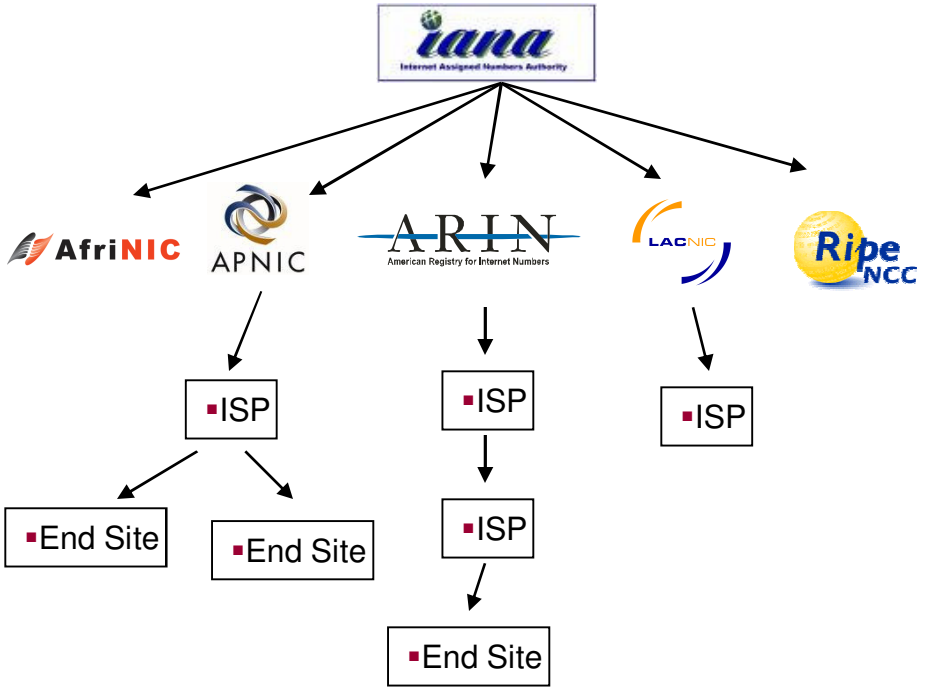

- **Internet Number Resource Allocation Hierarchy (AS numbers and IP addresses)**

IPv6 Allocations RIRs to LIRs/ISPs

How many allocations have been made by each RIR by year?

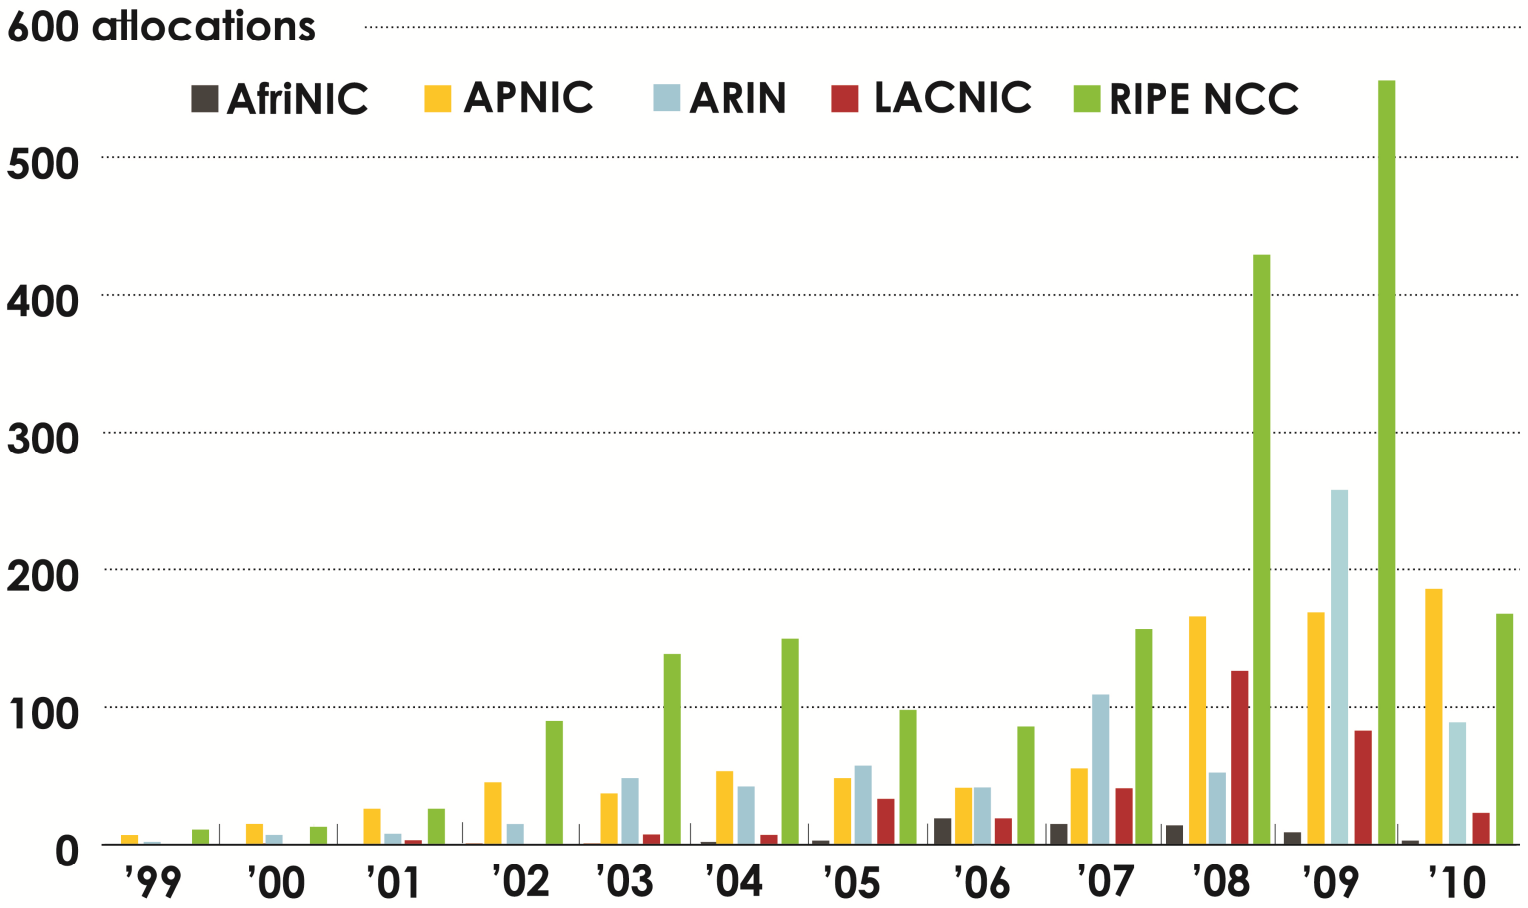

IPv4 & IPv6 Allocation Rate

Source: http://www.potaroo.net/tools/ipv4/index.html as of May 27, 2010

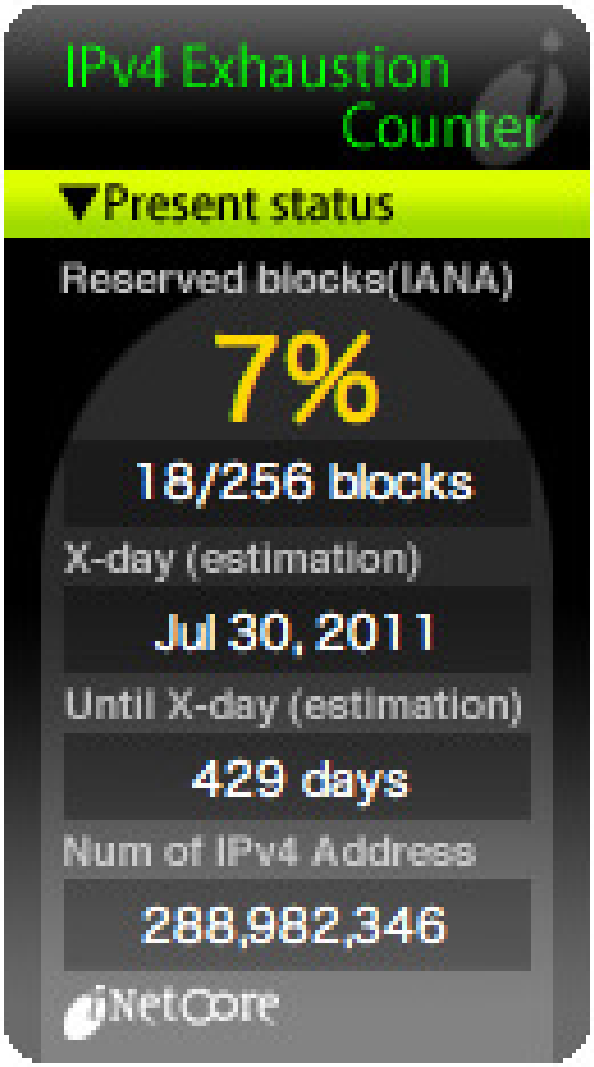

- Projected IANA unallocated IPv4 address pool exhaustion: 30-Jul-2011
- Projected RIR unallocated address pool exhaustion: 13-Mar-2012
- Those predictions based on current rates
- Projected support for IPv6 capabilities in your network: TBD!

Conceptual IP Router Architecture

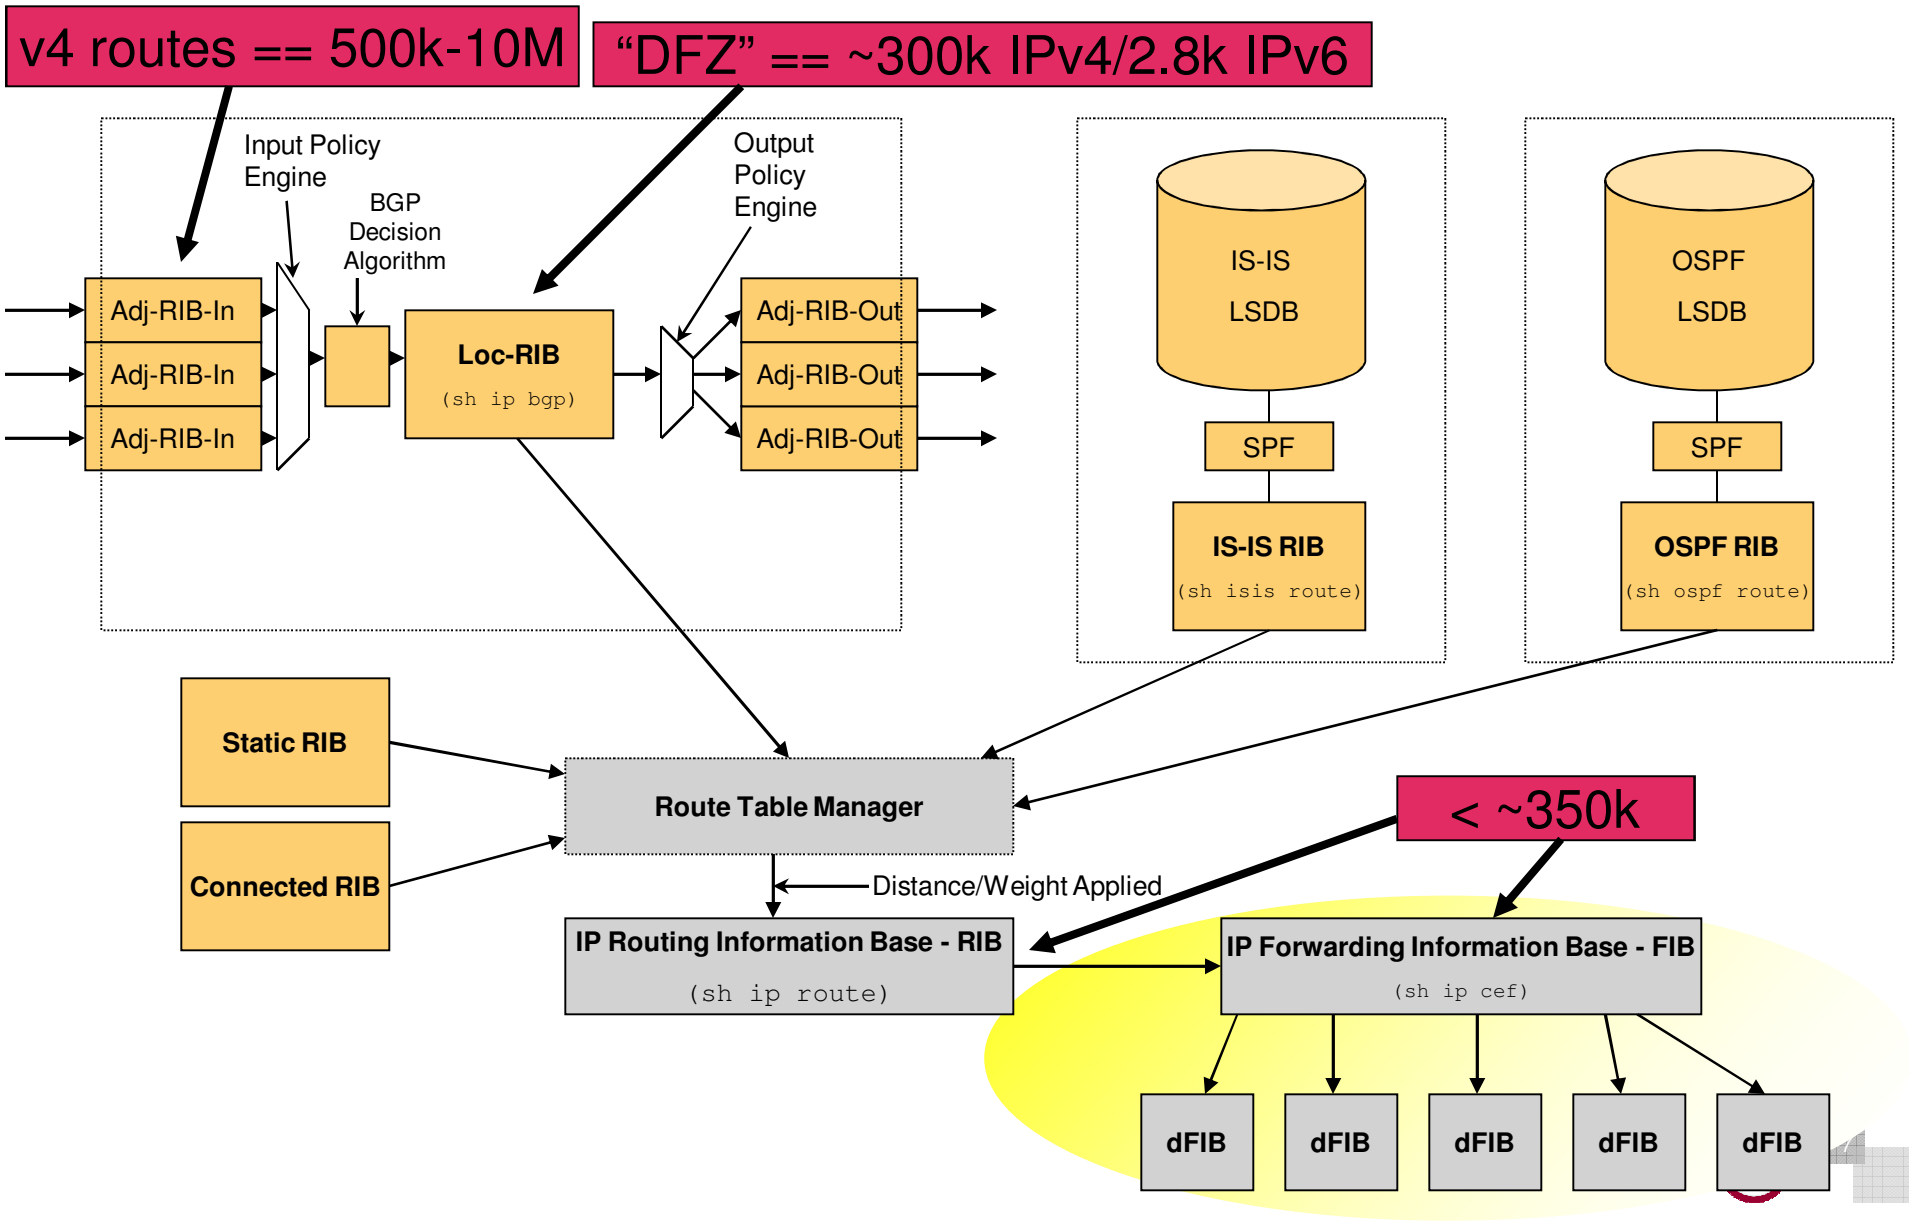

Internet Routing System Entries

Source: http://bgp.potaroo.net/v6/as2.0/index.html as of May 26, 2010

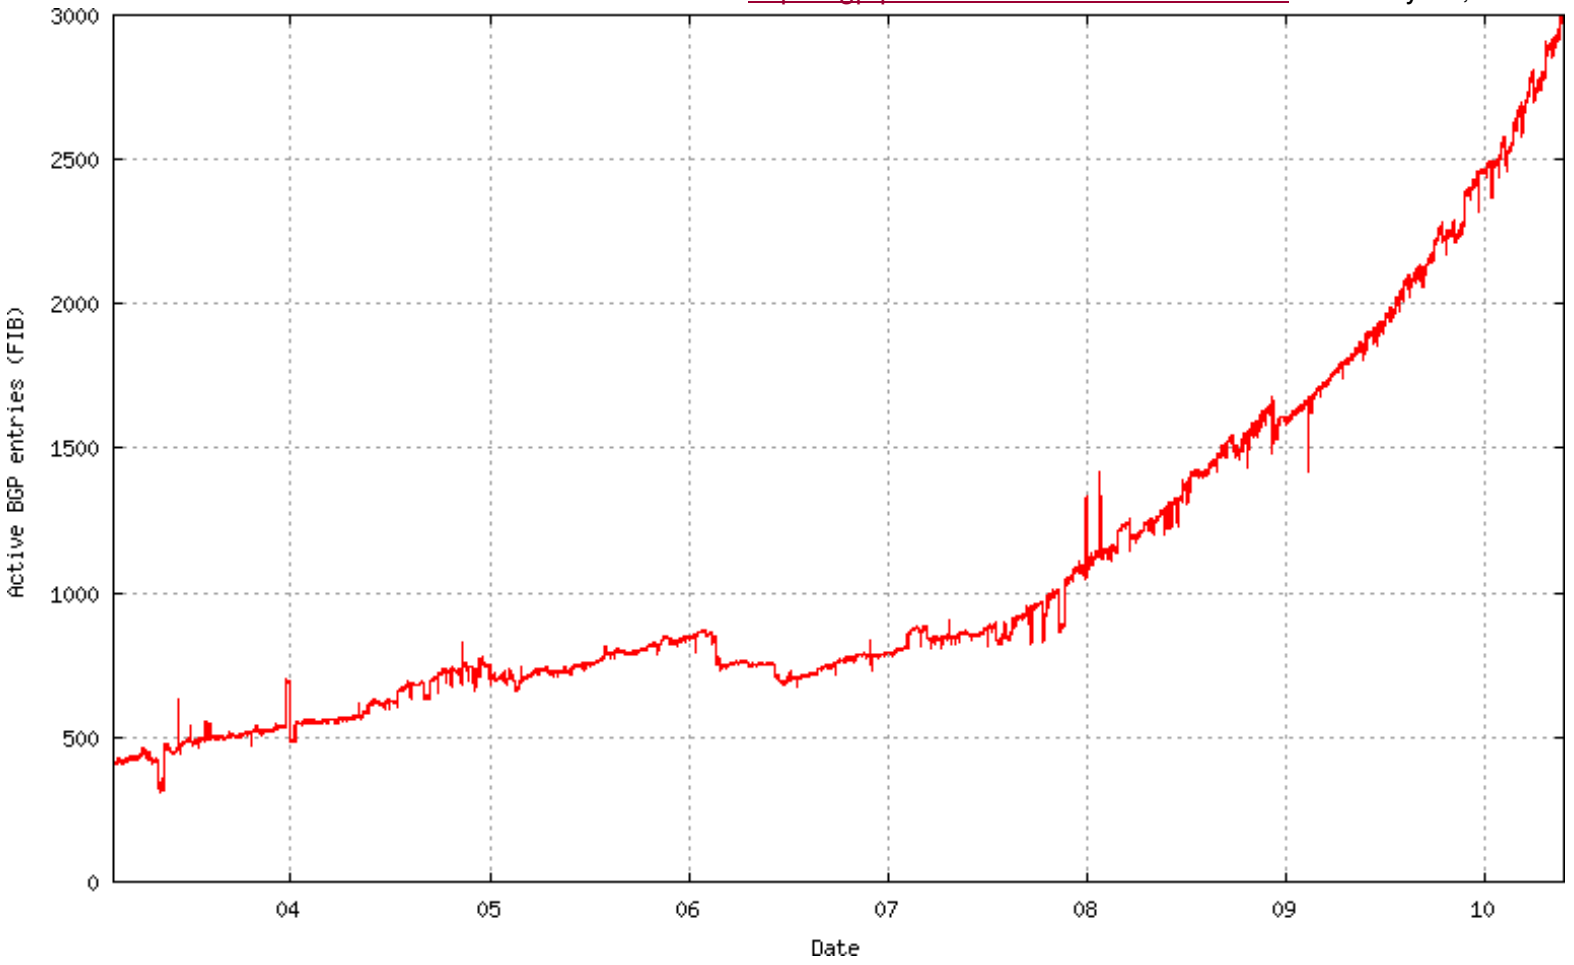

- 2,972 *Active* BGP RIB Entries
- Total *routes* dependent on location in topology

Domain Name Provisioning Functions

- Registrant selects registrar, interfaces with them
- Registrar interfaces with Registry for TLD
- Only one authoritative Registry per ccTLD/gTLD
- Some ccTLDs employ integrated registrar/registry model (e.g., .br)
- Does your registrar even support IPv6 NS glue record input?

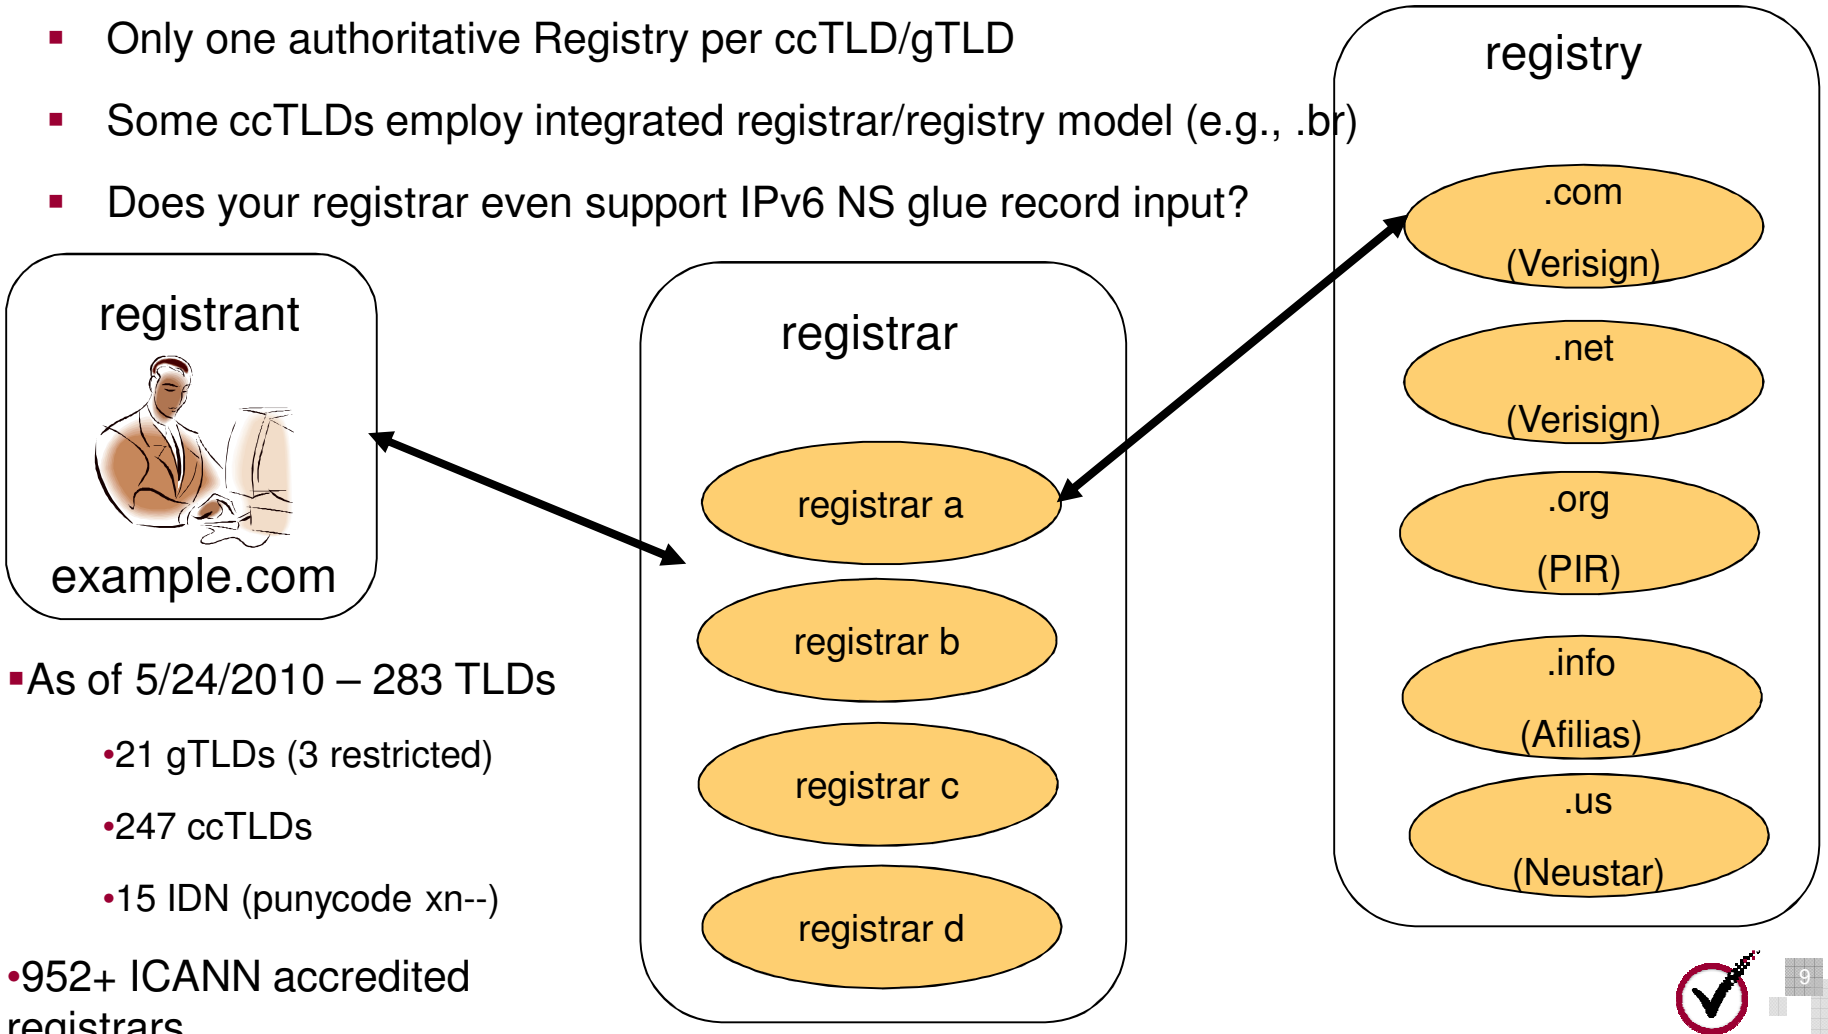

- As of 5/24/2010 – 283 TLDs
  - 21 gTLDs (3 restricted)
  - 247 ccTLDs
  - 15 IDN (punycode xn--)
- 952+ ICANN accredited registrars

Root Name Servers With IPv6 Records

- 13 root Internet name servers, operated by nearly as many discrete organizations
  - [a-m].root-servers.net
  - Only 7 currently have published IPv6 service addresses (AAAA records in the root zone)
    - A.ROOT-SERVERS.NET. 3600000 AAAA 2001:503:BA3E::2:30
    - F.ROOT-SERVERS.NET. 3600000 AAAA 2001:500:2F::F
    - H.ROOT-SERVERS.NET. 3600000 AAAA 2001:500:1::803F:235
    - J.ROOT-SERVERS.NET. 3600000 AAAA 2001:503:C27::2:30
    - K.ROOT-SERVERS.NET. 3600000 AAAA 2001:7FD::1
    - L.ROOT-SERVERS.NET. 3600000 AAAA 2001:500:3::42
    - M.ROOT-SERVERS.NET. 3600000 AAAA 2001:DC3::35
- Some of these folks use discrete IPv6 only infrastructure, as opposed to dual-stack systems

TLDs with IPv6 Records

Source: http://bgp.he.net/ipv6-progress-report.cgi as of May 25, 2010

- Criteria he.net folks apply:
  - To support IPv6 these nameservers should:
    - Have an IPv6 address themselves and native IPv6 connectivity so that they can be reached over IPv6.
    - Have AAAA records for their IPv6 address (glue records) in the root zone.
    - Be able to return AAAA (IPv6) address records.
  - Generated by parsing the root zonefile to get a list of TLDs and their associated nameservers, querying for AAAA records for the names of those nameservers and checking for AAAA records in the root zone for those nameservers.
- TLDs: 283
  - TLDs with IPv6 NS records: 228
  - Percentage of TLDs with IPv6 NS records: **80.6%**
  - TLDs with name servers with IPv6 glue in the root zone: 187
  - Percentage of TLDs that have name servers with IPv6 glue in the root zone: 66.1%

AAAA NS Glue in TLD Zone File

- As of May 24, 2010
  - .com – 578, .net – 860, .arpa – 10
  - Many more AAAAs provisioned, but large % inactive
- Full list available at he.net site
- Simply pulling zone files and querying for AAAA records without validation is misleading – LOTS of garbage out there, e.g.:

  danny@pork% dig bluecoat.com aaaa +short

  ::ffff:216.52.23.29
- http://tools.ietf.org/html/draft-morishita-dnsop-misbehavior-against-aaaa-00
- AAAA NS glue seems to be best measure of express enabled IPv6 sites at domain name level

Alexa 1M with IPv6 Records

Source: http://bgp.he.net/ipv6-progress-report.cgi as of May 25, 2010

- Top 1 million domains according to Alexa
  - 930,843 (93.1%) have resolvable ‘A’ record
  - 1,493 (0.15%) have resolvable ‘AAAA’ record
- Top 1 million domains with “www.” record
  - 972,325 (97.3%) have resolvable ‘A’ record
  - 2,076 (.21%) have resolvable ‘AAAA’ record
- Top 1 million domains with “ipv6.” record
  - 238,138 (23.8%) have resolvable ‘A’ record
  - 1,138 (.11%) have resolvable ‘AAAA’ record

Alexa 1M "Connectable"

Source: http://www.atoomnet.net/ipv6_enabled_popular_websites.php?complete_list=true as of May 25, 2010

- Popular sites (according to Alexa) with an IPv6 address on the main domain name (not subdomain)
  - Out of the 988840 tested websites only 2345 have one or more IPv6 addresses.
  - That is 0.24%. Out of the 4975 IPv6 addresses only 3478 are connectable.
  - Unfortunately, last updated 2/1/2010

TCP “failover” for IPv4 v. IPv6

- From a 24-hour snapshot of queries to VeriSign-operated TLDs
  - ~900kqps via IPv4 network layer
  - ~3kqps via IPv6 network layer
  - 25 v6 hosts exhibited failover to TCP, while 3213 IPv4 hosts exhibited failover
    - i.e., 2-3x the number of IPv6 hosts fail over to TCP v. IPv4 hosts
    - Of course, this is a crap statistic (need to compare failover / unique number of v4 or v6 hosts per day, as opposed to query rates)

- Working on this - lots more to come..

AMS-IX IPv6 Traffic Numbers

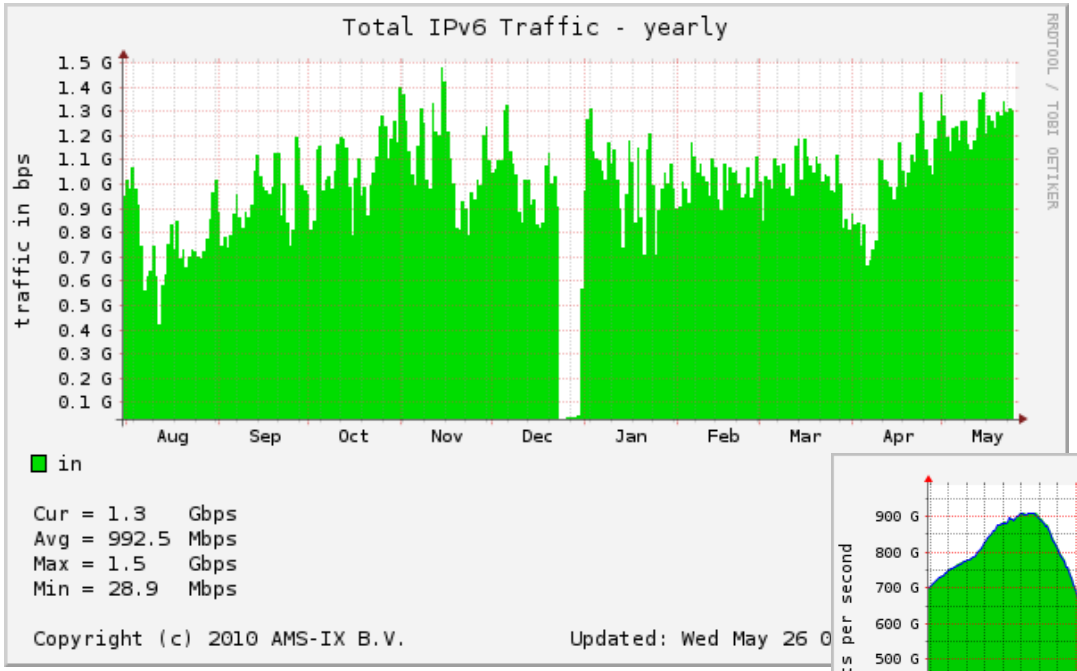

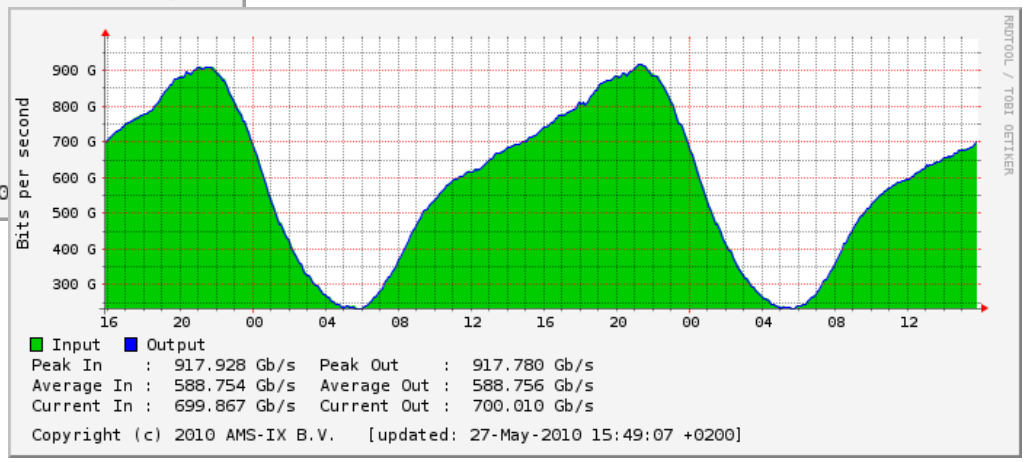

Source: http://www.ams-ix.net/sflow-stats/ipv6/

- Peak at about 900 Gbps
- IPv6 peaks about 1.5 Gbps (0.167%)

Arbor's Internet Observatory

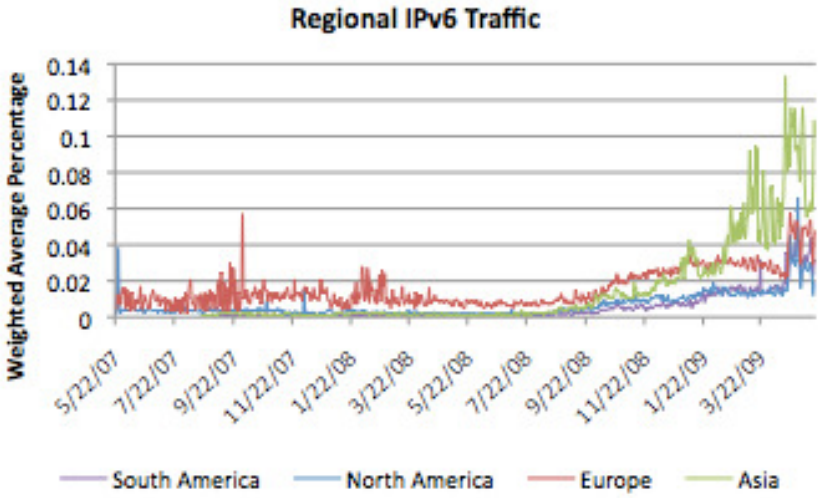

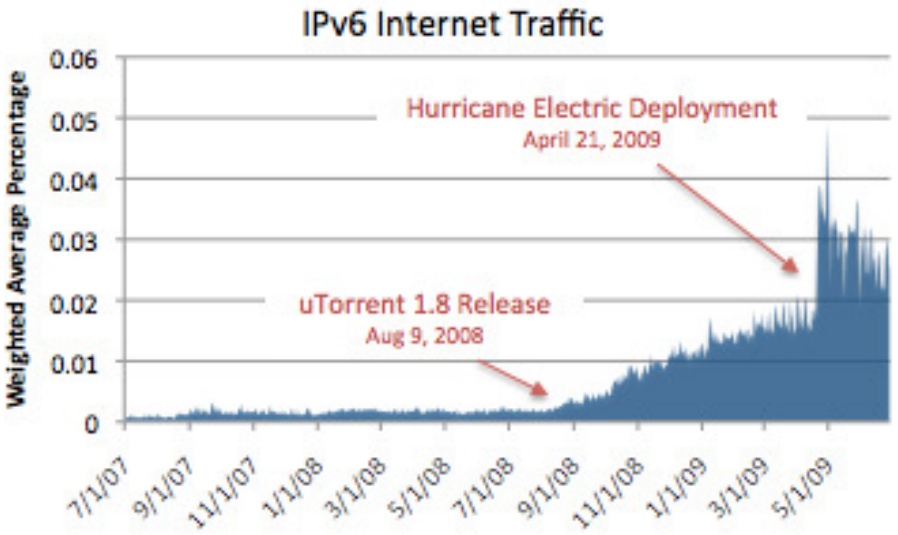

- August 2008 – IPv6 was less than 100th of 1%, IPv6 grew 1,400% in 12 months (mostly attributable to one application/ISP)
- Only 6 of 110 participants have native IPv6 router and collection infrastructure enabled (so this is just Teredo and 6to4 data)
- Only handful of participants use payload visibility, so only UDP Teredo control traffic is reported
- In the space of ten months uTorrent helped drive IPv6 traffic from .002% to .03% of all Internet traffic

Arbor Internet Observatory

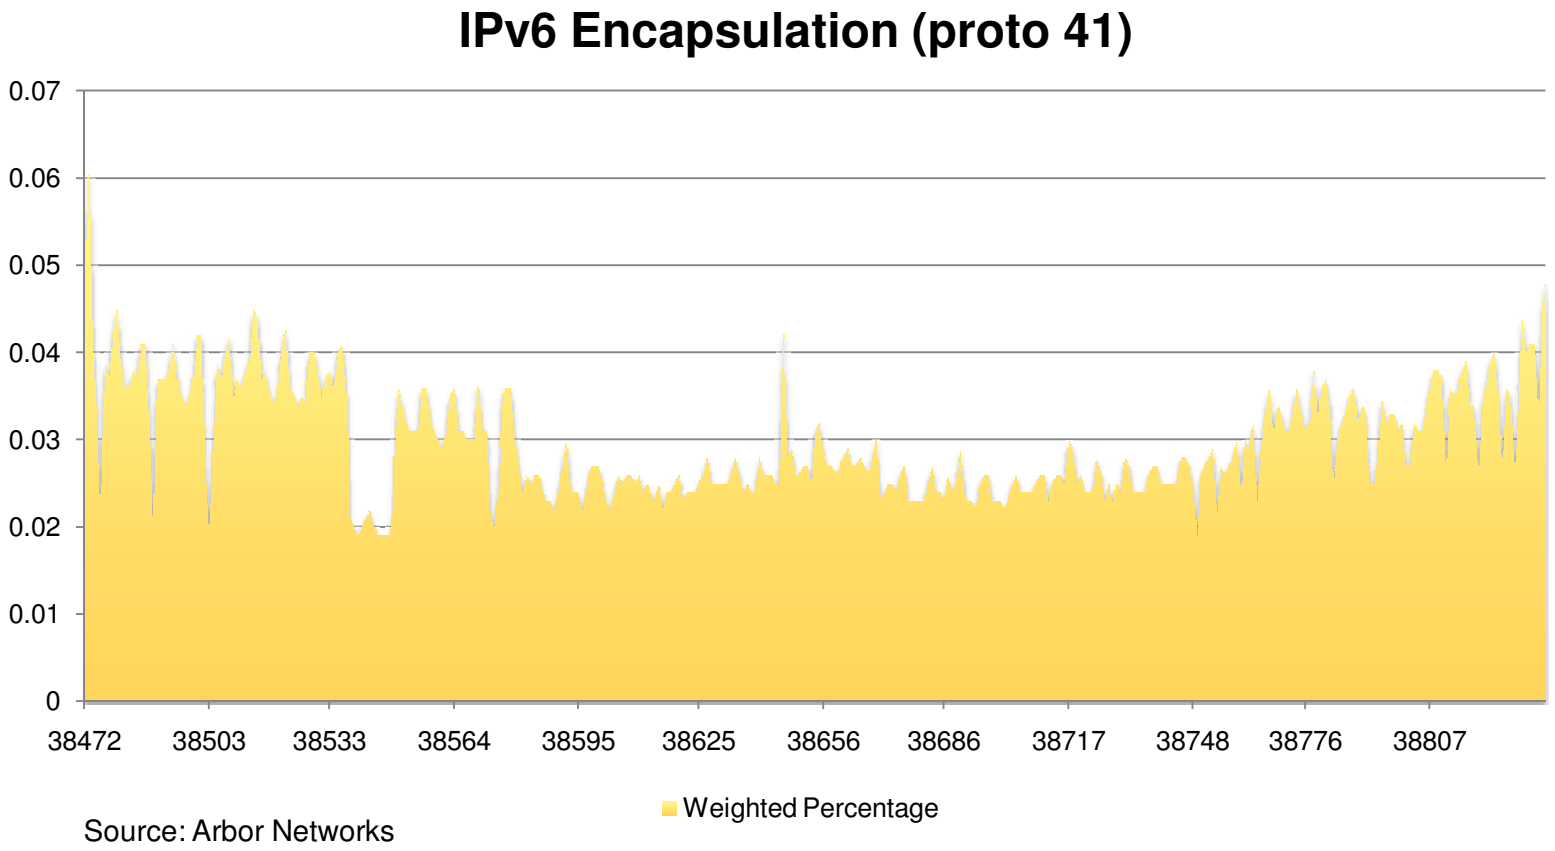

Source: Arbor Networks

- ~74 Internet Observatory participants in *this* dataset
- IP Proto 41 only != native IPv6
- These stats are just another data point – lack of native v6 visibility seriously limiting – reality is, we want *these* tunneled numbers to decrease

Arbor Internet Observatory

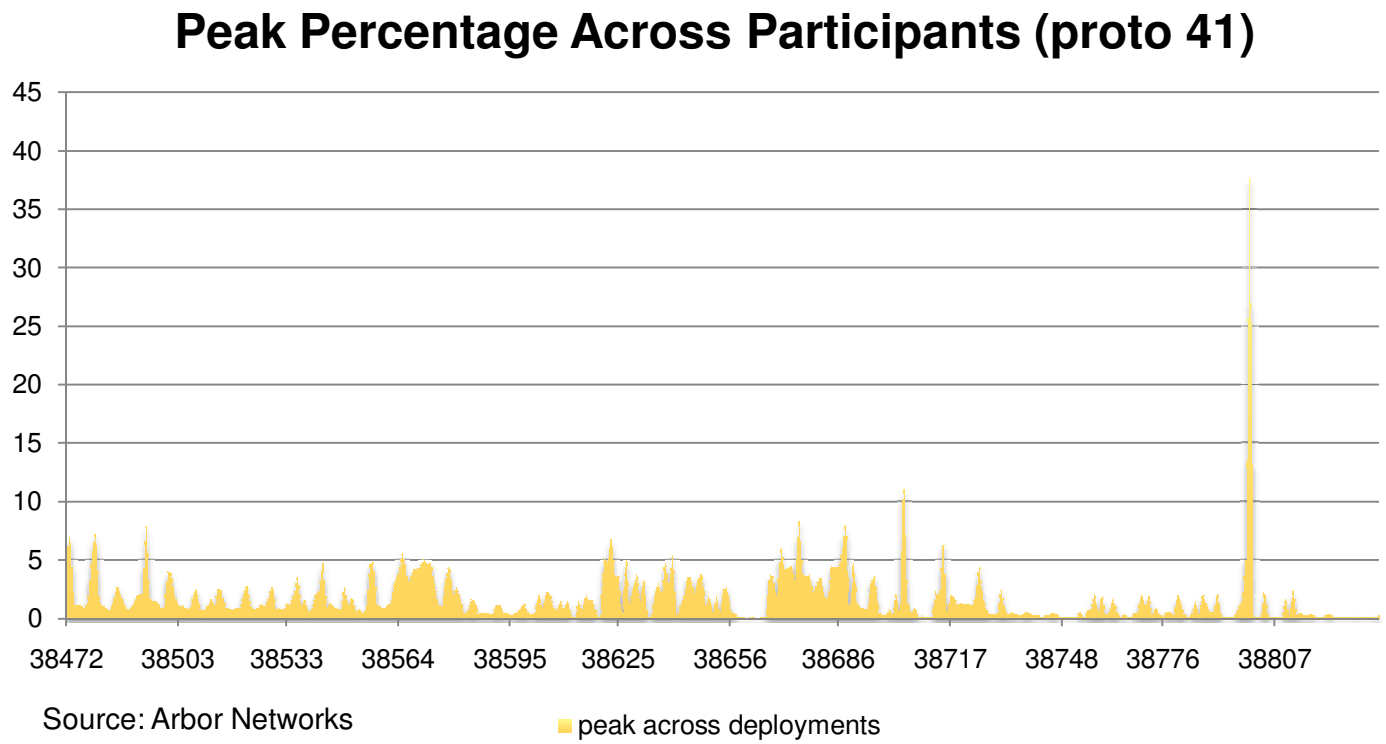

- Peak of 39.11% of single participants traffic on 3/25/2010

The End of End-to-End?

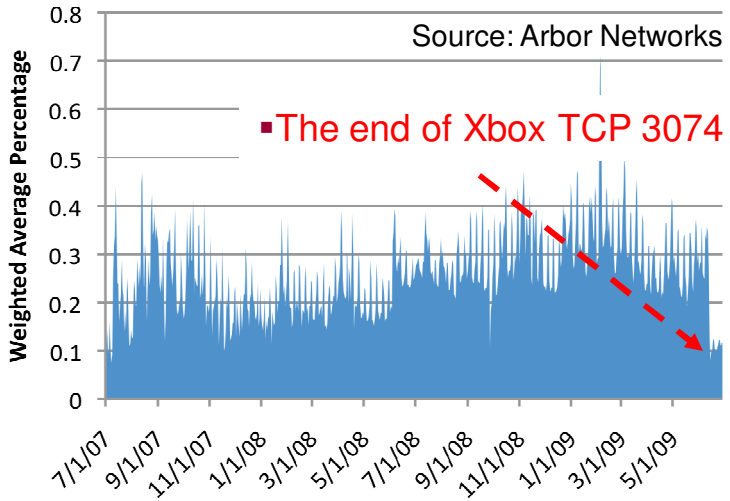

Weighted average percentage of Xbox Internet traffic

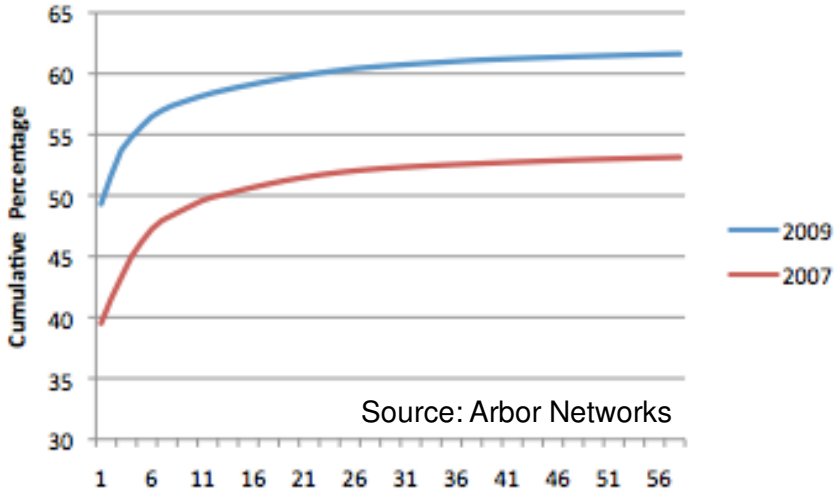

Cumulative Distribution of Traffic to TCP / UDP Ports

- Results in growing concentration of application traffic over a decreasing number of TCP / UDP ports
  - Especially port 80
  - Especially video
- Why this matters - little tolerance for brokenness, unnecessary latency, call center impact – IPv6 isn't going to make things any less complex initially – you need to be clued!

End of End Evolution…

- Today, application developers rather than being allocated discrete IP protocol numbers, or Transport layer ports, are developing their applications (and protocols) to work over IP/TCP/80 in order to operate effectively in most [firewalled] environments
- Expect IPv6 CG-NAT-PTesque devices to further compromise end-end transparency in the near future, will result in middleboxen for a very long time

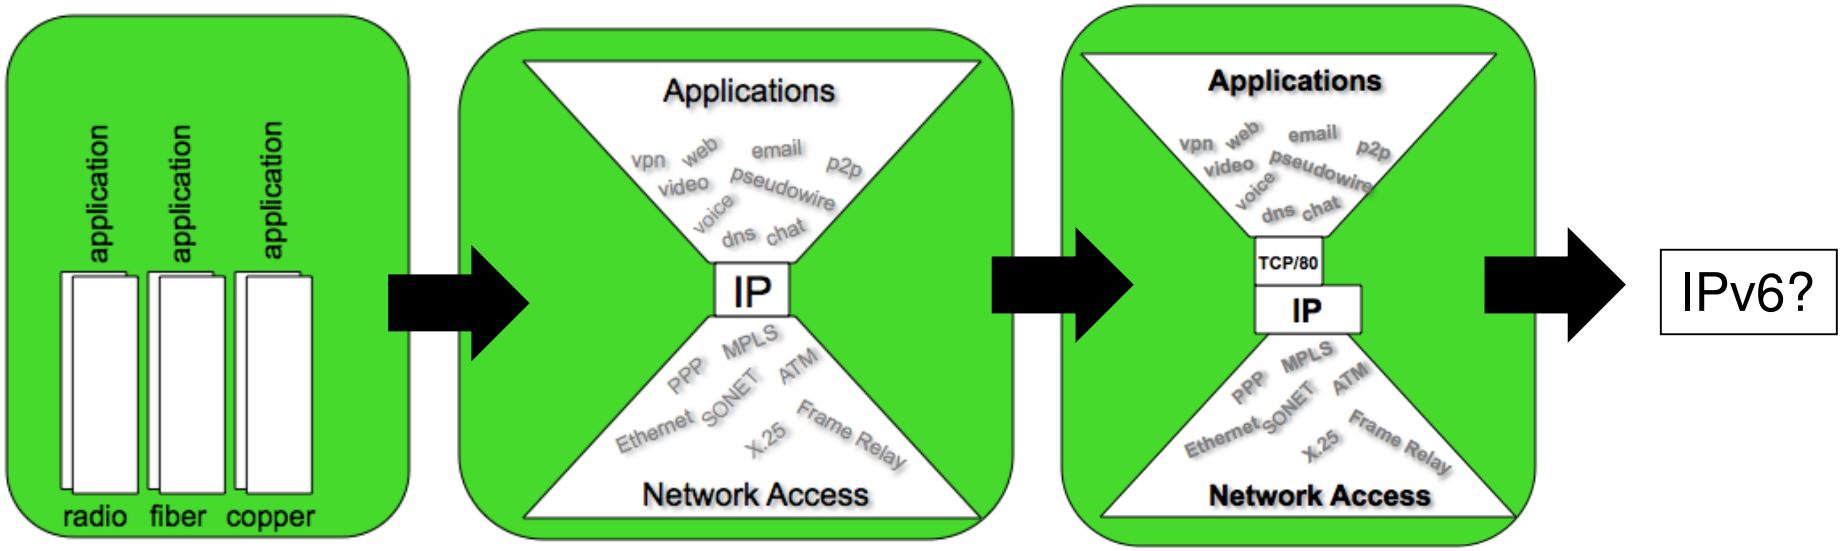

End System Capabilities

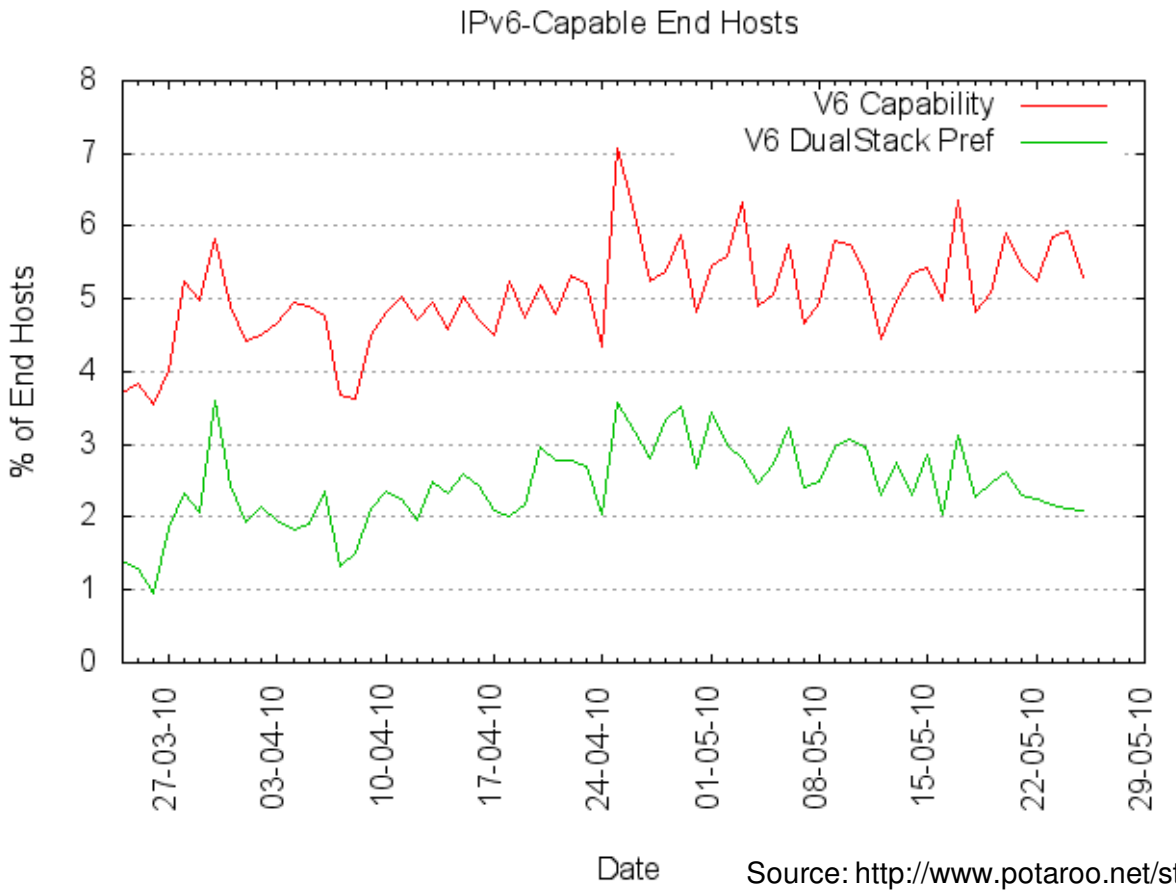

Source: http://www.potaroo.net/stats/1x1/

2% of the unique visitors to the APNIC web site are willing to prefer V6 when there is a choice, between 5% to 6% of end hosts will use V6 when they are not given the choice

OS Preferences

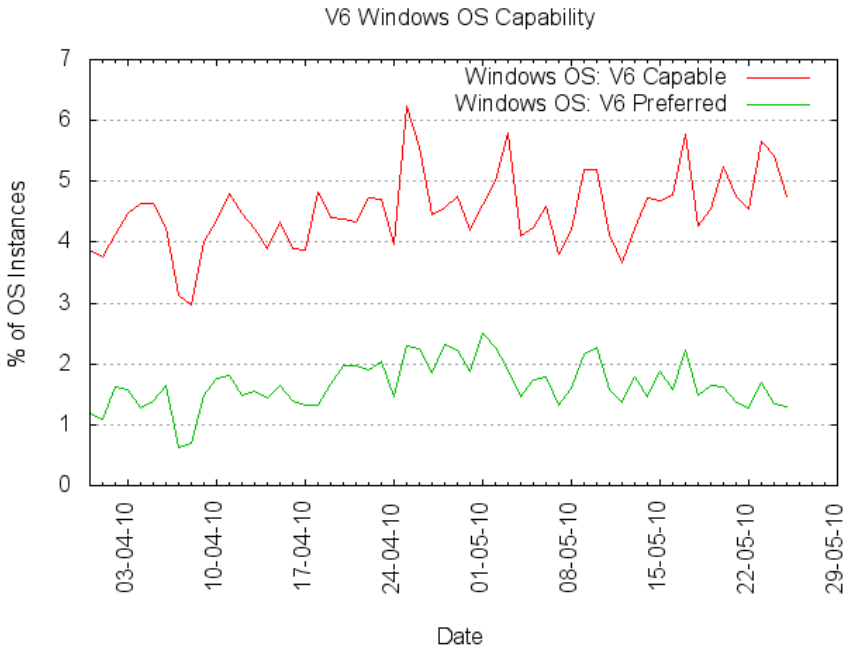

Source: http://www.potaroo.net/stats/1x1/

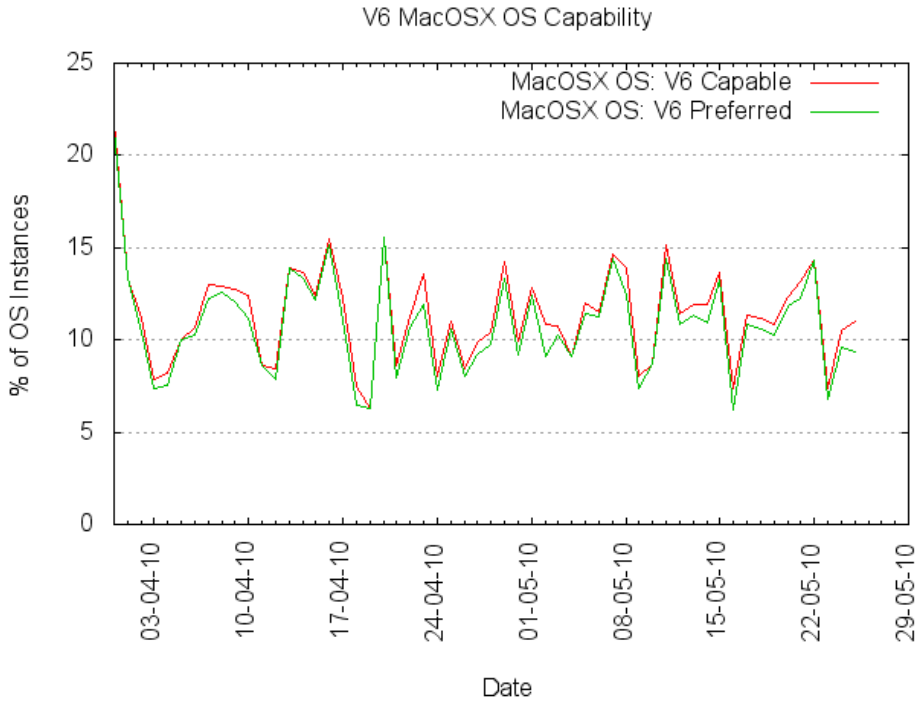

- These hosts that prefer to use IPv4 over IPv6 even in a dual stack choice are predominately Windows systems, MacOS systems strongly prefer IPv6. Reasoning appears to be choice of windows to ship with auto 6 over 4 tunneling mechanisms enabled with MacOS ships without such auto-tunnel mechanisms.

Browser Support & Layering Violations!

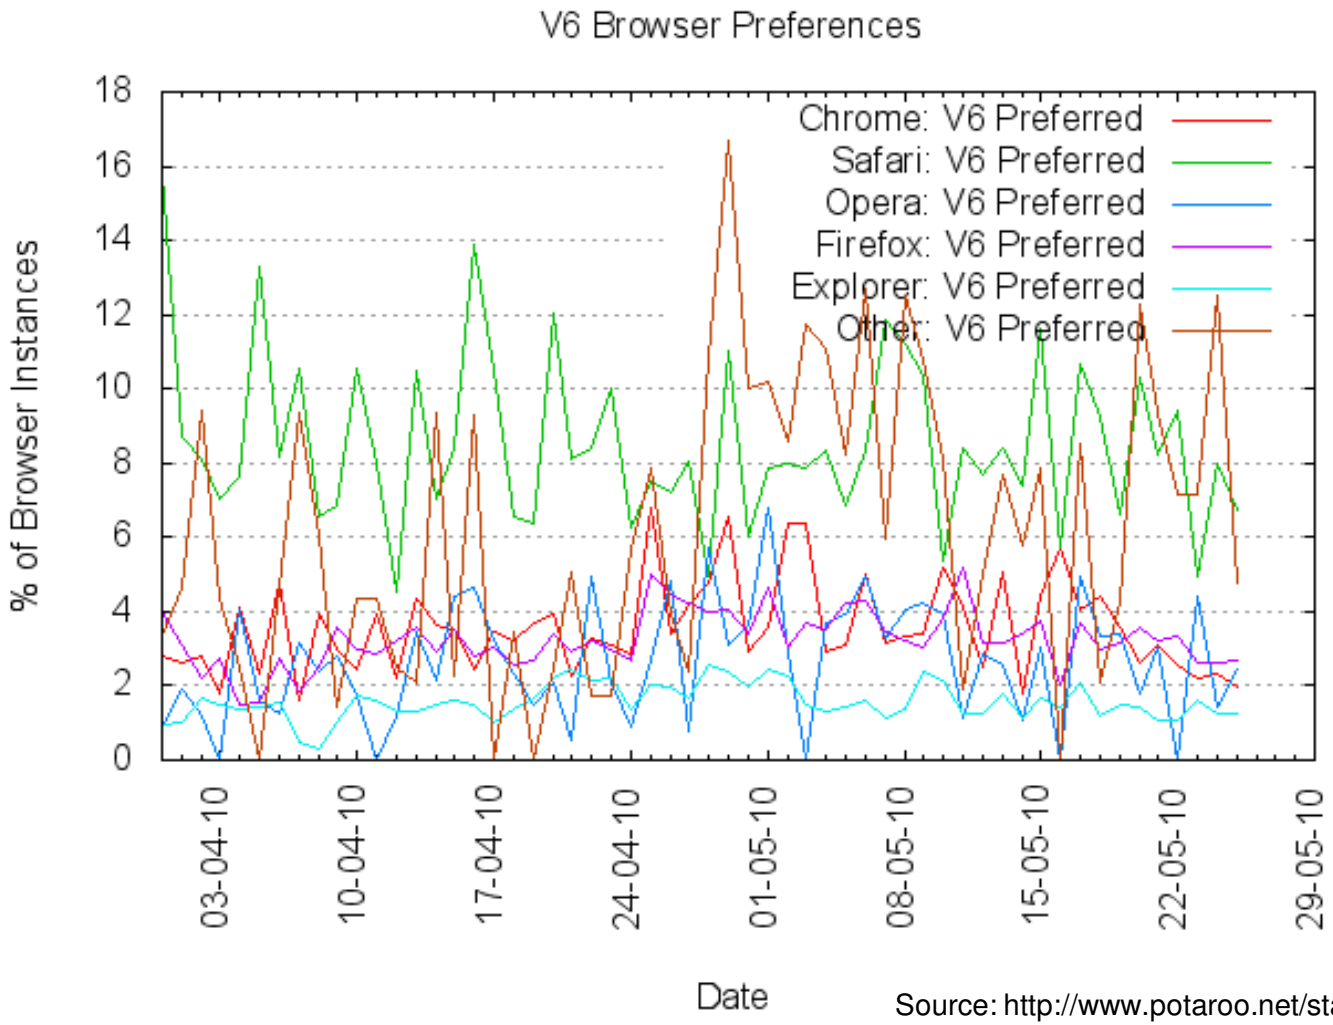

Source: http://www.potaroo.net/stats/1x1/

- IPv6 Dual Stack Preference for each browser type

RIPE Labs IPv6 Measurements Compilation

- http://labs.ripe.net/content/ipv6-measurements-compilation
  - APNIC
  - Comcast
  - Hurricane Electric
  - IPDN
  - Google
  - TNO & GNKS
  - Max.nl
  - Tore Anderson
  - Eric Vyncke
  - Mark Prior
  - CAIDA
  - RIPE NCC

Finally….

"Everybody's got a plan - until they get hit!"

--Mike Tyson

Summary

- A lot of the statistics are crap, a lot are very useful
- A lot of what you see as IPv6 is people measuring IPv6 (like bots)
- You can't manage what you don't measure
- IPv6 adoption brings systemic effects, you need to realize what the impacts of various components are in your operating environment
- Network layer visibility needs to be expanded
- End system and application layer instrumentation is best measure, verifies "service" functionality
- In a year, all bets are off the table!

Questions?

- *Statistics are like bikinis. What they reveal is suggestive, but what they conceal is vital. --- Aaron Levenstein*

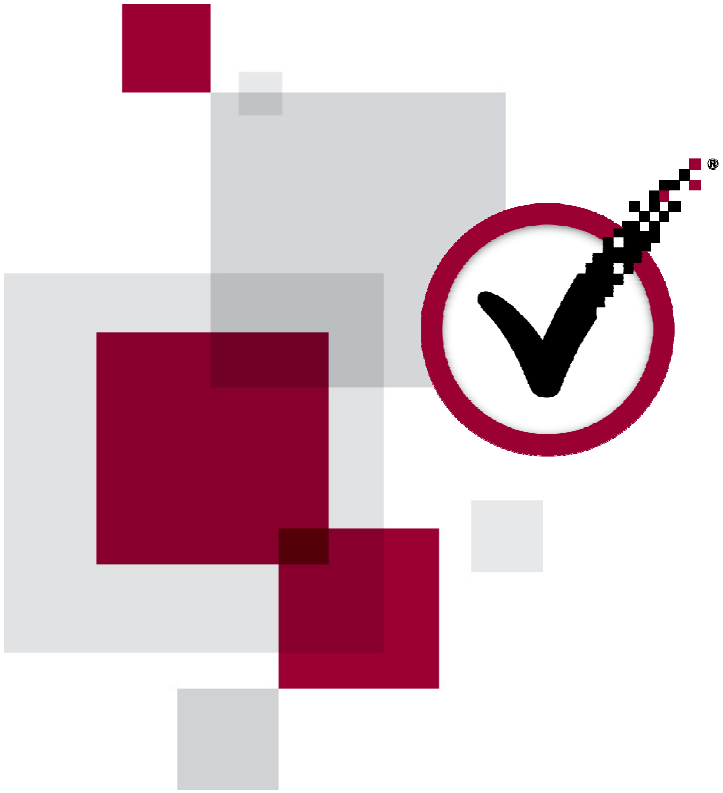

Thank You